arXiv:1608.01275v1 [cs.DS] 3 Aug 2016

The Dictionary Testing Problem

Siddharth Barman* Arnab Bhattacharyya* Suprovat Ghoshal*

Abstract

The *dictionary learning* (or *sparse coding*) problem plays an important role in signal processing, neuroscience, statistics and machine learning. Given a collection of vectors, the goal is to learn a basis with respect to which all the given vectors have a sparse representation. In this work, we study the testing analogue of this problem.

Roughly speaking, the *dictionary testing* problem is to decide whether a given collection of vectors can be approximately transformed into sparse vectors. More precisely, given as input vectors $\mathbf{y}_1, \ldots, \mathbf{y}_p \in \mathbb{R}^d$, we want to distinguish between the following two cases (for some choice of parameters):

(i) **YES case**: There exist a matrix $A \in \mathbb{R}^{d \times m}$ satisfying RIP and k-sparse unit vectors $\mathbf{x}_1, \ldots, \mathbf{x}_p \in \mathbb{R}^m$ such that $\mathbf{y}_i = A\mathbf{x}_i$ for all i.

(ii) **NO case**: There does not exist k'-sparse unit vectors $\mathbf{x}_1, \ldots, \mathbf{x}_p \in \mathbb{R}^m$ and $A \in \mathbb{R}^{d \times m}$ satisfying RIP such that $(A\mathbf{x}_1, \ldots, A\mathbf{x}_p)$ is "close" to $(\mathbf{y}_1, \ldots, \mathbf{y}_p)$.

Note that unlike known results in provable dictionary learning, there is no assumption on the distribution of $\mathbf{x}_1, \ldots, \mathbf{x}_p$ in the YES case. The goal is to design an efficient algorithm that can distinguish the above two cases using *linear queries*, i.e., inner products of the input vectors with vectors of our choice.

Our tester is very simple and efficient: it outputs **YES** iff the *gaussian width* of the input is sufficiently small. The analysis of the test can be viewed as giving a robust characterization of gaussian width in terms of sparsity. The number of linear queries made by the tester is *independent* of the dimension.

1 Introduction

Property testing is the study of algorithms that query their input a small number of times and distinguish between whether their input satisfies a given property or is "far" from satisfying that property. The quest for efficient testing algorithms was initiated by [BLR93, BFL91] and later explicitly formulated by [RS96, GGR98]. Property testing can be viewed as a relaxation of the traditional notion of a decision problem, where the relaxation is quantified in terms of a distance parameter. There has been extensive work in this area over the last couple of decades; see, for instance, the surveys [Ron08, RS06] for some different perspectives.

As described by Ron in her survey [Ron08], property testing can also alternatively be viewed as a relaxation of machine learning. In property testing, one wishes to quickly check whether a hypothesis about an input object is approximately true, whereas in learning, one wants to obtain an approximate representation of the input object assuming a hypothesis holds. It is easy to see that (proper) learning assuming a hypothesis implies testing of that hypothesis. Also, the connection to learning gives a natural motivation for property testing; the tester can be a preliminary step taken before an expensive learning algorithm is run, in order to check whether to use a particular class of functions as the hypothesis class for the learner.

The vast majority of work in property testing has dealt with properties of *discrete* objects, such as bipartiteness of graphs or linearity of functions $f : \mathbb{F}_2^n \to \mathbb{F}_2$ or membership in error-correcting codes or representability of functions as concise boolean formulae. However, these properties are very unlike the hypothesis classes used in modern machine learning applications, such as in computer vision, speech recognition and feature extraction. In these settings, the input data often consists of real vectors and the accuracy parameter is in terms of

*Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India. Emails: {barman, arnabb, suprovat.ghoshal}@csa.iisc.ernet.in.

certain norms of real vector spaces. Hence, it is natural to ask if we can develop efficient testers for interesting properties of *continuous* objects (such as real vectors and matrices) that are either more efficient or have more provable guarantees than the counterpart learning algorithms?

Motivated by these considerations, we examine the testing analog of the *dictionary learning* or *sparse coding* problem. The goal of the dictionary learning problem is to express a collection of input vectors as linear combinations of a small number of vectors chosen from a large *dictionary*. Dictionary learning is a fundamental task in several domains. The problem was first formulated by Olshausen and Field [OF96, OF97] who showed that the dictionary elements learnt from sparse coding of natural images have certain similarities to the receptive fields of neurons in the visual cortex. Neuroscientists have also extracted dictionaries for speech [LS00] and video [Ols02] inputs. Inspired by these results, automatically learned dictionaries have been used in machine learning for feature selection [EP07] and for denoising [EA06], edge-detection [MLB+08], super-resolution [YWHM08], restoration [MSE08], and texture synthesis [Pey09] in image processing applications. Dictionary learning is also a component in some deep learning systems [RBL07].

Given a matrix $Y \in \mathbb{R}^{d\times p}$, the decomposition to be learnt by a dictionary learning algorithm is:

$$Y = AX, \qquad A \in \mathbb{R}^{d\times m}, X \in \mathbb{R}^{m\times p} \tag{1}$$

where each column of X has at most k non-zero entries. The columns of Y correspond to the input samples, the columns of A correspond to the dictionary atoms, and the columns of X are the mixing coefficients for each sample. The dictionary size m may be much larger than d, the so-called "overcomplete" setting. In the usual setup, the problem is studied using a *generative model* in which the columns of X are sampled iid from some fixed distribution on sparse vectors.

The subject of this work is the *dictionary testing* problem, in which the goal is to decide whether a decomposition of the form in (1) exists or whether Y is "far" from any such decomposition. Clearly, the testing problem is a relaxation of dictionary learning; the tester need not recover A and X. Indeed, a motivation behind our work was devising a fast tester that can quickly rule out inputs not admitting a sparse coding. In this paper, we establish provable guarantees for a dictionary testing algorithm, even when the columns of X are arbitrary sparse vectors (normalized to have unit norm). Guarantees in this *agnostic* setting are much stronger than in the generative models which are typically used to analyze dictionary learning algorithms.

1.1 Problem Setup

Let $\mathcal{S}^{d-1}$ be the unit sphere in d dimensions. We denote by $\mathsf{Sp}_k^d \subset \mathcal{S}^{d-1}$ the set of all k sparse unit vectors, i.e., $\mathsf{Sp}_k^d = \{\mathbf{x} : \mathbf{x} \in \mathcal{S}^{d-1}, \|\mathbf{x}\|_0 \leqslant k\}$. We drop the superscript d on Sp_k^d whenever it is clear from context.

We will devise a tester that approximately decides if the input can be sparsely coded by a dictionary satisfying the *restricted isometry property (RIP)*.

Definition 1.1 ((ε, k)-Restricted Isometry Property)**.** *For $\varepsilon > 0$ and $S \subseteq \mathbb{R}^m$, a matrix $A \in \mathbb{R}^{d\times m}$ is said to be ε-*isometric on S *if:*

$$(1-\varepsilon)\|\mathbf{x}-\mathbf{y}\|_2^2 \leqslant \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leqslant (1+\varepsilon)\|\mathbf{x}-\mathbf{y}\|_2^2$$

*for all $\mathbf{x}, \mathbf{y} \in S$. A matrix $A \in \mathbb{R}^{d\times m}$ is said to be (ε, k)-*RIP *if it is ε-isometric on Sp_k^m.*

RIP is a very natural assumption to put on dictionary matrices. The seminal work of Candes, Romberg and Tao [CRT06] showed that given a matrix A satisfying $(0.4, k)$-RIP, any $\mathbf{x} \in \mathsf{Sp}_k^m$ can be recovered efficiently from $A\mathbf{x}$, even when the sparsity $k = \Omega(d)$. Moreover, many natural families of redundant dictionaries, such as those constructed from wavelets and sinusoids, or from wavelets and ridgelets, satisfy the RIP criterion [DH01]. Indeed, these dictionaries are *incoherent*; a dictionary $A \in \mathbb{R}^{d\times m}$ is μ-incoherent if the inner product between any two columns is at most $\mu/\sqrt{d}$. Incoherence is a strictly stronger restriction[1] than RIP and was identified in [GMS03] as underlying many overcomplete dictionaries that are used widely by domain experts.

[1] See, for example, Lemma A.7 of [AAJ+14] for a formal statement and proof.

In the dictionary learning context, [AAJ+14] proved that dictionaries satisfying RIP can be learnt using an alternating minimization algorithm.

To decide when input vectors are "far" from being sparse, we use the following measure:

Definition 1.2. *Given sets $S, S' \subset \mathbb{R}^d$ and $\varepsilon \geqslant 0$, we say that S is* ε-isoangular *to S' if there exists a bijection $\psi : S \to S'$ such that for all $u, v \in S$,*

$$\left| \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\| \cdot \|\mathbf{v}\|} - \frac{\langle \psi(\mathbf{u}), \psi(\mathbf{v}) \rangle}{\|\psi(\mathbf{u})\| \cdot \|\psi(\mathbf{v})\|} \right| \leqslant \varepsilon.$$

In words, S is ε-isoangular to S' if the pairwise angles between vectors in S are approximately equal to the pairwise angles between corresponding vectors in S', upto an additive error ε. Note that, for a set of vectors S and matrix A, where S and $A(S)$ are all (approximately) unit norm, then if A is ε-isometric, then S is $O(\varepsilon)$-isoangular with $A(S)$. The notion of ε-isoangular is weaker than ε-isometry in three senses. First of all, the notion of ε-isoangular is invariant to scaling whereas ε-isometry clearly is not. Secondly, Definition 1.2 gives a bound on the additive error between angles, while the error bound in Definition 1.1 is multiplicative. Thirdly, the map ψ is required to be linear in Definition 1.1.

We assume that the input vectors are accessed using *linear queries*. This means that for an input vector $\mathbf{y} \in \mathbb{R}^d$, the algorithm can obtain $\langle \mathbf{v}, \mathbf{y} \rangle$ where $\mathbf{v}$ is any vector in $\mathbb{R}^d$ using one linear query. These are exactly the queries allowed in the setting of compressed sensing, dimensionality reduction using random linear projections, and in many differentially private mechanisms [HT10, BDKT12, HKR12]. We define the *query complexity* of a dictionary tester to be the number of linear queries made for the worst possible input.

1.2 Our Results

Our main result is the construction and analysis of an efficient algorithm for the dictionary testing problem, that accepts vectors sparsely represented by a dictionary satisfying RIP and rejects inputs that are not almost isoangular to vectors sparsely represented by a dictionary satisfying RIP.

Theorem 1.3 (Testing sparse representations)**.** *Fix $\varepsilon, \delta \in (0,1)$ and positive integers d, k, m and p, such that $(k/d)^{1/8} < \varepsilon < \frac{1}{100}$, $k \geqslant 4$, and $m \geqslant 20k\varepsilon^{-4}$. There exists a randomized testing algorithm which makes linear queries to input vectors $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_p \in \mathbb{R}^d$ and has the following properties (where Y is the matrix having $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_p$ as columns):*

- **Completeness***: If $Y = AX$ where $A \in \mathbb{R}^{d \times m}$ satisfies (ε, k)-RIP and $X \in \mathbb{R}^{m \times p}$ with each column of X in Sp_k^m, then the tester accepts with probability $\geqslant 1 - \delta$.*
- **Soundness***: If there does not exist a matrix $A \in \mathbb{R}^{d \times m}$ satisfying (ε, k)-RIP and vectors $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p \in \mathsf{Sp}_\ell^m$ such that $\{\mathbf{y}_i : i \in [p]\}$ is $O(\varepsilon^{1/8})$-isoangular to $\{A\mathbf{x}_i : i \in [p]\}$, then the tester rejects with probability $\geqslant 1 - \delta$. Here, $\ell = O(k/\varepsilon^4)$.*

The query complexity of the tester is $O(p\varepsilon^{-2} \log(p/\delta))$.

The soundness guarantee roughly says that the tester will reject when there is no (ε, k)-RIP dictionary of size m such that the input vectors are nearly isoangular to linear combinations of $\ell = O_\varepsilon(k)$ of the dictionary elements. We again emphasize that the result holds in the *agnostic* setting, without any distributional assumption on the inputs.

An easy application of the same techniques used to prove Theorem 1.3 also yields a tester for dimensionality that may be of independent interest.

Theorem 1.4 (Testing dimensionality)**.** *There exists a randomized testing algoirthm which takes as input $\varepsilon, \delta \in (0,1)$, positive integers d, k and p, and unit vectors $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_p \in \mathcal{S}^{d-1}$. It makes linear queries to $\mathbf{y}_1, \ldots, \mathbf{y}_p$ and has the following properties:*

- **Completeness***: If $\dim(\mathbf{y}_1, \ldots, \mathbf{y}_p) \leqslant k$, then the tester accepts with probability $\geqslant 1 - \delta$.*

– **Soundness**: *If there exist no* $\mathbf{y}'_1, \ldots, \mathbf{y}'_p$ *of dimension* $\leqslant k'$ *that are* ε*-isometric to* $\mathbf{y}_1, \ldots, \mathbf{y}_p$*, then the tester rejects with probability at least* $1 - \delta$*. Here,* $k' = O(k/\varepsilon^2)$*.*

The query complexity of the tester is $O(p \log(\delta^{-1}))$*.*

Note that here, the soundness guarantee is stronger i.e., the tester rejects when the input is not almost isometric to a low-dimensional subspace, unlike Theorem 1.3 where the soundness guarantee is in terms of isoangularity.

1.3 Discussion

For dictionary testing, we are intuitively looking for a quantity ω that *robustly* captures sparsity and is easily computable using linear queries, in the sense that ω is small when the input vectors are sparse and large when they are "far" from sparse. Moreover, ω needs to be invariant with respect to isometries and nearly invariant with respect to near-isometries.

A natural and widely-used measure of structure that is approximately invariant to isometries is the *gaussian width*.

Definition 1.5. *The* gaussian width *of a set* $S \subseteq \mathbb{R}^d$ *is:*

$$\omega(S) = \mathop{\mathbf{E}}_{\mathbf{g}}[\sup_{\mathbf{v} \in S} \langle \mathbf{g}, \mathbf{v} \rangle]$$

where $\mathbf{g} \in \mathbb{R}^d$ *is a random vector drawn from* $N(0,1)^d$*, i.e., a vector of independent standard normal variables.*

The gaussian width of S measures how well on average the vectors in S correlate with a randomly chosen direction. It is invariant under orthogonal transformations of S as the distribution of $\mathbf{g}$ is spherically symmetric. The following bounds on the gaussian width are well-known.

Lemma 1.6 (See, for example, [RV08, Ver15])**.** *Below,* C *denotes a universal constant, which could be different each time it appears:*

(i) If S *is a finite subset of* $\mathcal{S}^{d-1}$*, then* $\omega(S) \leqslant C\sqrt{\log |S|}$*.*

(ii) $\omega(\mathcal{S}^{d-1}) \leqslant \sqrt{d}$*.*

(iii) If $S \subseteq \mathcal{S}^{d-1}$ *is of dimension* k*, then* $\omega(S) \leqslant \sqrt{k}$*.*

(iv) $\omega(\mathsf{Sp}_k^d) \leqslant 2\sqrt{3k \log(d/k)}$ *when* $d/k > 2$ *and* $k \geqslant 4$*.*

We take advantage of the last two items in Lemma 1.6 to prove the completeness direction of Theorem 1.4 and Theorem 1.3. The gaussian width of a set S can be easily estimated because Gaussian concentration of measure implies that with high probability, the random variable $\sup_{\mathbf{y} \in S} \langle \mathbf{g}, \mathbf{y} \rangle$ is close to its expectation $\omega(S)$. Hence, it is natural to test dimensionality or sparsity by computing an empirical estimate of the gaussian width and comparing the estimate to the results in Lemma 1.6.

The challenge then is to show that when the input is far from having the desired structure, the gaussian width also shows deviation. Contrapositively, we need to prove that small gaussian width implies closeness to a particular structure. For this, we use the fact that the square of the gaussian width, $\omega^2(S)$, can be interpreted as the *effective dimension* of S. Unlike the linear algebraic notion of dimension, gaussian width is robust to perturbations such as random projections. This connection (formalized by Gordon's Theorem, see Theorem 2.1) immediately proves Theorem 1.4: if the gaussian width of the input set of vectors is $\ll \sqrt{k}$, then they can be nearly isometrically projected to a space of dimension $\ll k/\varepsilon^2$.

For Theorem 1.3, the argument is more involved. Again, by Gordon's Theorem, if the gaussian width of the input set S is $\ll \sqrt{k \log(d/k)}$, then S can be ε-isometrically projected into a space of dimension $d' \ll k\varepsilon^{-2} \log(d/k)$. Next, we show that these projected vectors are close (in terms of isoangularity) to a random projection of $\mathsf{Sp}^m_{k\varepsilon^{-4}}$ into d' dimensions. In this part of the argument, which is the technical heart of this paper, we actually want the random projection to *distort* $\mathsf{Sp}^d_{k\varepsilon^{-4}}$ so that after normalization, the projection of the sparse vectors forms a net of $\mathcal{S}^{d'-1}$.

1.4 Related Work

Although the dictionary learning problem is NP-hard [DMA97] in the worst case (even when the dictionary A is known), it is generally considered a solved problem in practice. Gradient descent heuristics [OF97], the method of optimal directions (MOD) [EAH99] and the K-SVD algorithm [AEB05] work very well for real applications; see [Aha06] for a very nice overview. On the other hand, theoretical results have not yet been able to fully justify the efficacy of these methods. The first work to give a dictionary learning algorithm with provable guarantees was [SWW12] who restricted the dictionary to be square and the sparsity to be at most $\sqrt{d}$. For the more common overcomplete setting, Arora et al. [AGM14] and Agarwal et al. [AAJ$^+$14] independently gave algorithms with provable guarantees for dictionaries satisfying incoherence and RIP respectively. These works also restrict the sparsity to be strictly less than $\sqrt{d}/\mu$ where μ is the incoherence. Barak et al. [BKS15] gave a very different analysis using the sum-of-squares hierarchy that works for nearly linear sparsity; however, their algorithm runs in time $d^{\mathrm{poly}(1/\varepsilon)}$ where ε measures the accuracy to which the dictionary is to be learned and this is too inefficient to be of use for realistic parameter ranges. All of these (as well as other more recent) works assume distributions from which the input samples are generated in an iid fashion.

In contrast, our work is in the agnostic setting. We do not require the number of samples p to be large, as in [AAJ$^+$14, AGM14, BKS15], while our guarantees hold even when the sparsity k is as large as $m/\log m$. Moreover, since our tester performs a very simple empirical averaging to estimate the gaussian width of the input, it is much faster than any of the learning algorithms.

We are not aware of any work in the testing literature that studies the sparse coding problem. Czumaj et al. [CSZ00] studied some problems in computational geometry from the property testing perspective, but the problems involved only discrete structures. Krauthgamer and Sasson [KS03] studied the problem of testing dimensionality, but their notion of farness from being low-dimensional is quite different from ours. In their setup, a sequence of vectors $\mathbf{y}_1, \ldots, \mathbf{y}_p$ is ε-far from being d-dimensional if at least εp vectors need to be removed to make it be of dimension d. Note that a set of vectors can be nearly isometric to a d-dimensional subspace but far from being d-dimensional in Krauthgamer and Sasson's sense (for example, the Johnson-Lindenstrauss projection of the standard unit vectors $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_d$).

2 Preliminaries

We use the following version of Gordon's Theorem repeatedly in this work.

Theorem 2.1 (Gordon's Theorem [Gor85]). *Given $S \subset \mathcal{S}^{D-1}$ and $\mathbf{G} \sim \frac{1}{\sqrt{d'}} N(0,1)^{d' \times D}$, we have*

$$\mathop{\mathbf{E}}_{\mathbf{G}} \left[\max_{\mathbf{x} \in S} \|\mathbf{G}\mathbf{x}\|_2 \right] \leqslant 1 + \frac{\omega(S)}{\sqrt{d'}}$$

It directly implies the following generalization of the Johnson-Lindenstrauss lemma.

Theorem 2.2 (Generalized Johnson-Lindenstrauss lemma). *Let $S \subseteq \mathbb{R}^n$. Then there exists linear transformation $\Phi : \mathbb{R}^n \mapsto \mathbb{R}^{d'}$, for $d' = O\left(\frac{\omega(S)^2}{\varepsilon^2}\right)$, such that Φ is an ε-isometry on S. Moreover, $\Phi \sim \frac{1}{\sqrt{d'}} N(0,1)^{d' \times n}$ is an ε-isometry on S with high probability.*

It can be easily verified that the quantity $\max_{\mathbf{x} \in S} \|\mathbf{G}\mathbf{x}\|_2$ is 1-Lipschitz. Therefore, using Gaussian concentration for Lipschitz functions, we get the following corollary :

Corollary 2.3. *Let S and G be as in Theorem 2.1. Then for all $\varepsilon > 0$, we have*

$$\Pr_{\mathbf{G}} \left(\max_{\mathbf{x} \in S} \|\mathbf{G}\mathbf{x}\|_2 \geqslant 1 + (1+\varepsilon)\frac{\omega(S)}{\sqrt{d'}} \right) \leqslant \exp\left(- O(\varepsilon \omega(S))^2 \right)$$

The following lemma gives concentration for the gaussian width:

Lemma 2.4 (Concentration on the gaussian width [BLM13]). *Let $S \subset \mathbb{R}^d$. Let $W = \sup_{\mathbf{v}\in S}\langle \mathbf{g}, \mathbf{v}\rangle$ where $\mathbf{g}$ is drawn from $N(0,1)^d$. Then:*

$$\mathbf{Pr}[|W - \mathbf{E}\, W| > u] < 2e^{-\frac{u^2}{2\sigma^2}}$$

where $\sigma^2 = \sup_{\mathbf{v}\in S}\left(\|\mathbf{v}\|_2^2\right)$. Notice that the bound is dimension independent.

We shall also need the following well-known comparison inequality relating suprema of gaussian processes:

Lemma 2.5 (Slepian's lemma [Sle62]). *Let $\{X_u\}_{u\in U}$ and $\{Y_u\}_{u\in U}$ be two almost surely bounded centered Gaussian processes, indexed by the same compact set U. If for every $u_1, u_2 \in U$:*

$$\mathbf{E}\left[|X_{u_1} - X_{u_2}|^2\right] \leqslant \mathbf{E}\left[|Y_{u_1} - Y_{u_2}|^2\right]$$

then we have

$$\mathbf{E}\left[\sup_{u\in U} X_u\right] \leqslant \mathbf{E}\left[\sup_{u\in U} Y_u\right]$$

2.1 Algorithmic Estimation of Gaussian Width and Norm of a vector

We record here simple lemmas bounding the number of linear queries needed to estimate the gaussian width of a set and the length of a vector.

Lemma 2.6 (Estimating Gaussian Width using linear queries). *For any $u > 4$, $\varepsilon \in (0, 1/2)$ and $\delta > 0$, there is a randomized algorithm that given a set $S \subseteq \mathbb{R}^d$ and $\|\mathbf{v}\| \in [1 \pm \varepsilon]$ for all $\mathbf{v} \in S$, computes $\hat{\omega}$ such that $\omega(S) - u \leqslant \hat{\omega} \leqslant \omega(S) + u$ with probability at least $1 - \delta$. The algorithm makes $O(\log(1/\delta) \cdot |S|)$ linear queries to S.*

Proof. By Lemma 2.4, for a random $\mathbf{g} \sim N(0,1)^d$, $\sup_{\mathbf{v}\in S}\langle \mathbf{g}, \mathbf{v}\rangle$ is away from $\omega(S)$ by u with probability at most $2e^{-16/4.5} < 0.1$. By the Chernoff bound, the median of $O(\log \delta^{-1})$ trials will satisfy the conditions required of $\hat{\omega}$ with probability at least $1 - \delta$. □

Lemma 2.7 (Estimating norm using linear queries). *Given $\varepsilon \in (0, 1/2)$ and $\delta > 0$, for any vector $\mathbf{x} \in \mathbb{R}^d$, only $O(\varepsilon^{-2} \log \delta^{-1})$ linear queries to $\mathbf{x}$ suffice to decide whether $\|\mathbf{x}\| \in [1 - \varepsilon, 1 + \varepsilon]$ with success probability $1 - \delta$.*

Proof. It is easy to verify that $\mathbf{E}_{\mathbf{g}\sim N(0,1)^d}[\langle \mathbf{g}, \mathbf{x}\rangle^2] = \|\mathbf{x}\|^2$. Therefore, it can be estimated to a multiplicative error of $(1 \pm \varepsilon/2)$ by taking the average of the squares of linear measurements using $O\left(\frac{1}{\varepsilon^2}\log\frac{1}{\delta}\right)$-queries. For the case $\|\mathbf{x}\|_2 \leqslant 2$, a multiplicative error $(1 \pm \varepsilon/2)$ implies an additive error of ε. Furthermore, when $\|\mathbf{x}\|_2 \geqslant 2$, a multiplicative error of $(1 \pm \varepsilon/2)$ implies that $L \geqslant 2(1 - \varepsilon/2) > 1 + \varepsilon$ for $\varepsilon < 1/2$. □

2.2 Analysis of Dimensionality Tester

As an illustration of the techniques involved in proving Theorem 1.3, we give a simple proof of Theorem 1.4 here.

Proof of Theorem 1.4. Let S denote the set $\{\mathbf{y}_1, \ldots, \mathbf{y}_p\}$. The tester obtains $\hat{\omega}$ that approximates $\omega(S)$ to an additive error of $\sqrt{k}$ and accepts iff $\hat{\omega} \leqslant 2\sqrt{k}$. By Lemma 2.6, the tester requires $O(p \log \delta^{-1})$ linear queries to obtain $\hat{\omega}$.

If $\dim(S) \leqslant k$, then by Lemma 1.6, $\hat{\omega} \leqslant 2\sqrt{k}$ with probability at least $1 - \delta$, so that the tester accepts with the same probability.

If the tester accepts, then with probability at least $1 - \delta$, $\omega(S) \leqslant 3\sqrt{k}$. By Theorem 2.2, S is then ε-isometric to a subspace of dimension $O(k/\varepsilon^2)$. □

3 Analysis of Dictionary Tester

In the rest of this paper, we prove Theorem 1.3. Let S denote the set $\{\mathbf{y}_1, \ldots, \mathbf{y}_p\}$. Our testing algorithm is as follows:

Algorithm 1: SparseTestAgnostic

1 Use Lemma 2.7 to decide with probability at least $1 - \delta/2$ if there exists $\mathbf{y}_i$ such that $\|\mathbf{y}_i\| \notin [1 - 2\varepsilon, 1 + 2\varepsilon]$. Reject if so.
2 Use Lemma 2.6 to obtain $\hat{\omega}$, an estimate of $\omega(S)$ within additive error $\sqrt{3k \log(m/k)}$ with probability at least $1 - \delta/2$.
3 Accept if $\hat{\omega} \leqslant 4\sqrt{3k \log(m/k)}$, else reject

The number of linear queries made by the tester is $O(p\varepsilon^{-2} \log(p/\delta))$ in Line 1 and $O(p \log \delta^{-1})$ in Line 2.

3.1 Completeness

Assume that for each $i \in [p]$, $\mathbf{y}_i = A\mathbf{x}_i$ for a matrix $A \in \mathbb{R}^{d \times m}$ satisfying (ε, k)-RIP and $\mathbf{x}_i \in \mathsf{Sp}_k^m$. By definition of RIP, we know that $1 - \varepsilon \leqslant \|\mathbf{y}_i\| \leqslant 1 + \varepsilon$, so that Line 1 of the algorithm will pass with probability at least $1 - \delta/2$.

From Lemma 1.6, we know that $\omega(\{\mathbf{x}_1, \ldots \mathbf{x}_p\}) \leqslant 2\sqrt{3k \log(m/k)}$. Lemma 3.1 shows that the gaussian width of S is approximately the same; its proof, deferred to the appendix, uses Slepian's Lemma (Lemma 2.5).

Lemma 3.1. *Let $X \subset \mathcal{S}^{m-1}$ be a finite set, and let $S \subset \mathbb{R}^d$ be an ε-isometric embedding of X. Then*

$$(1 - \varepsilon)\omega(X) \leqslant \omega(S) \leqslant (1 + \varepsilon)\omega(X) \tag{2}$$

The gaussian width of $\mathbf{y}_1, \ldots, \mathbf{y}_p$ is $\leqslant 2(1 + \varepsilon)\sqrt{3k \log(m/k)}$. Taking into account the additive error in Line 2, we see that with probability at least $1 - \delta/2$, $\hat{\omega} \leqslant (3 + 2\varepsilon)\sqrt{3k \log(m/k)} \leqslant 4\sqrt{3k \log(m/k)}$. Hence, the tester accepts with probability at least $1 - \delta$.

3.2 Soundness

As mentioned before, in order to prove soundness we need to show that whenever the gaussian width of the set S is small, it is *close* to some sparse point-set. The following lemma, which in a sense shows an approximate inverse to Gordon's Theorem for sparse vectors, makes the statement precise:

Lemma 3.2 (Inverse Gordon)**.** *Given $S \subset \mathcal{S}^{d-1}$ such that $\omega(S) = O(\sqrt{k \log(m/k)})$. Then, for $\ell = O(k/\varepsilon^4)$, there exists a matrix $\Phi \in \mathbb{R}^{d \times m}$ satisfying (ε, k)-RIP and $\tilde{S} \subset \mathsf{Sp}_\ell^m$ such that $\Phi(\tilde{S})$ is $O(\varepsilon^{1/8})$-isoangular to S.*

Given this lemma, we can finish the soundness analysis. Let S' denote the set where each element $\mathbf{y}_i$ of S is replaced by the normalized vector $\mathbf{y}_i/\|\mathbf{y}_i\|_2$. If Line 1 of the tester passes, then $\omega(S') \leqslant (1 + 3\varepsilon)\omega(S)$. Now, if Line 2 of the tester also passes, then invoking Lemma 3.2 on S', we obtain that there exists a (ε, k)-RIP matrix $\Phi \in \mathbb{R}^{d \times m}$ and $\tilde{S} \subset \mathsf{Sp}_\ell^m$ such that $\Phi(\tilde{S})$ is $O(\varepsilon^{1/8})$-isoangular to S', and hence to S.

4 Proof of Inverse Gordon Lemma

Our proof uses the probabilistic method. We show that with positive probability, there exists a matrix Φ satisfying RIP such that Φ of a sparse point-set is nearly isoangular to the target set S. At a very high level, the proof can be broken down into two parts:

Quantity	Interpretation
d	Native dimension of the input
d'	Lower dimension as given by the JL lemma
ε	User input for error tolerance
Φ	Scaled Gaussian matrix as given by the JL lemma of dimension $d' \times d$
Φ^{norm}	Normalized analogue of Φ
ϕ_i	i^{th} Block of Φ of dimensions $d' \times Nd'$
k	Sparsity parameter for completeness
ℓ	Sparsity parameter for soundness
$\mathcal{S}^{d-1}$	Unit sphere in d-dimensions
Sp_k^m	Set of k-sparse points in $\mathbb{R}^m$
$\widehat{\mathsf{Sp}}_\ell^m$	Discretized set of ℓ-sparse points in $\mathbb{R}^m$

Figure 1: Table of Notations

- **Projection**: Sets with small Gaussian width point sets can be almost isometrically embedded into a low dimensional subspace.
- **Covering**: Appropriately sparse point sets on projection to a low dimensional subspace form a cover of the unit sphere on the smaller dimension.

The first point directly follows from the Generalized Johnson-Lindenstrauss Lemma (Theorem 2.2). The second point is formalized as follows:

Definition 4.1 (Discretized sparse vectors)**.**

$$\widehat{\mathsf{Sp}}_\ell^m = \left\{ \mathbf{x} : \mathbf{x} \in \left\{ 0, \pm\frac{1}{\sqrt{\ell}} \right\}^m, \|\mathbf{x}\|_0 = \ell \right\}$$

Theorem 4.2. *Given $\varepsilon > 0$ and integers C, d, k and m, let $d' = \frac{Ck}{\varepsilon^2}\log(m/k)$. Suppose $m \gg k/\varepsilon^4$ and $d \gg d'$. Let $\Phi : \mathbb{R}^m \mapsto \mathbb{R}^{d'}$ be drawn from $\frac{1}{\sqrt{d'}}N(0,1)^{d' \times m}$. Then, for $\ell = \frac{2Ck}{\varepsilon^4}$, $\Phi^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell^m)$ is a $20\varepsilon^{1/4}$-cover of $\mathcal{S}^{d'-1}$, where $\Phi^{\text{norm}}(\mathbf{x}) = \Phi(\mathbf{x})/\|\Phi(\mathbf{x})\|_2$.*

Assuming Theorem 4.2, let us finish the proof of Lemma 3.2.

Proof of Lemma 3.2. We assume that $\mathbf{0} \in S$, since this does not change the gaussian width of the set. Let $\Phi_1 \in \mathbb{R}^{d' \times d}$ be as given by the generalized JL lemma (Theorem 2.2), so that $\Phi_1(S)$ is ε-isometric to S. Note that due to isometry, $\Phi_1(S)$ is ε-close to $\mathcal{S}^{d'-1}$ in the ℓ_2-distance.

Let $\Phi \in \mathbb{R}^{d' \times m}$ be as in Theorem 4.2. By Theorem 2.2, Φ satisfies (ε, k)-RIP w.h.p. Furthermore, since $\Phi^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ is an $O(\varepsilon^{1/4})$ cover of $\mathcal{S}^{d'-1}$, it is also an $O(\varepsilon^{1/4})$ cover of $\Phi_1(S)$. Therefore, there exists $\tilde{S} \subset \widehat{\mathsf{Sp}}_\ell$ such that $|\tilde{S}| = |S|$, and

$$\forall \mathbf{x} \in \Phi_1(S) \quad \exists \mathbf{p}_\mathbf{x} \in \Phi^{\text{norm}}(\tilde{S}) \text{ s.t. } \|\mathbf{x} - \mathbf{p}_\mathbf{x}\|_2 \leqslant C_3\varepsilon^{1/4} \tag{3}$$

Now $\Phi_1(S)$ is an ε-isometric embedding of S in d' dimensions. Furthermore $\Phi^{\text{norm}}(\tilde{S})$ is an $O(\varepsilon^{1/4})$ cover of $\Phi_1(S)$ (from Eq. 3). Therefore from Proposition A.5 we have $\Phi^{\text{norm}}(\tilde{S})$ is $O(\varepsilon^{1/8})$-isoangular to S in $\mathbb{R}^d$. Hence, we get that $\Phi(\tilde{S})$s is $O(\varepsilon^{1/8})$-isoangular to S as well, since isoangularity is invariant to scaling of the vectors. Finally, observe that $\Phi \in \mathbb{R}^{d' \times m}$. Since $d' \leqslant d$, and therefore we can pad Φ with $d - d'$ rows of $\mathbf{0}$, which would give us a linear operator of dimensions $d \times m$ with identical guarantees, hence establishing the lemma.

□

4.1 Proof of Theorem 4.2

The intent here is to get the sparse point set $\widehat{\mathsf{Sp}}_\ell$ *distorted* on projection, so that it forms an ε-cover of the unit sphere on the smaller dimension. However, doing so rules out proofs that rely on simple union bound arguments. For instance, on allowing the projections to become distorted, we run into the risk of lots of points collapsing together into a small fraction of $\mathcal{S}^{d'-1}$. As a result, the set $\Phi^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ could turn out to be insufficient for forming a cover of the unit sphere. These issues are avoided by carefully relating the gaussian width of $\Phi^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ to that of $\mathcal{S}^{d'-1}$, followed by a partitioning argument. The partitioning crucially uses the block structure of elements in $\widehat{\mathsf{Sp}}_\ell$, which results in independent and distributionally identical blocks, allowing us to take union bounds effectively.

Proof. The proof of this Theorem proceeds in steps. We first partition Φ into L blocks of Nd' columns each, for some appropriately chosen N. So, $\Phi = [\phi_1 \cdots \phi_L]$. Note that ϕ_i is a $d' \times Nd'$ submatrix of $\Phi \in \mathbb{R}^{d' \times m}$. Write $\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ to denote $\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell^{Nd'})$. Also, note that $\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ is equal to the set obtained by applying Φ to vectors in $\widehat{\mathsf{Sp}}_\ell$ whose support is contained in P, where $P \subset [m]$ is the set of Nd' columns of Φ that are present in ϕ_i.

For any such fixed partition ϕ_i, we show that the restriction $\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ has large expected gaussian width (Lemma 4.3), where again, $\phi_i^{\text{norm}}(\mathbf{x}) = \phi_i(\mathbf{x})/\|\phi_i(\mathbf{x})\|$. Furthermore, using Lemma 4.4 we argue that any fixed point on $\mathcal{S}^{d'-1}$ has distance $O(\varepsilon^{1/4})$ to $\phi_i(\widehat{\mathsf{Sp}}_\ell)$ with large probability. Now, we use the independence of the ϕ_i's to argue that the probability of $\mathbf{x} \in \mathcal{S}^{d'-1}$ being simultaneously far away from all $\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ is exponentially small. Finally, taking a union bound over the ε-net of $\mathcal{S}^{d'-1}$ completes the proof.

We fix the parameter N to be

$$N = \sqrt{\frac{m}{k}}\Big(\log \frac{m}{k}\Big)^{-1} \tag{4}$$

Set the number of blocks $L = \frac{m}{Nd'} = O(\varepsilon^2\sqrt{m/k})$ blocks of Nd' coordinates each. By construction, for any fixed $i \in [L]$, $\phi_i \sim \frac{1}{\sqrt{d'}}N(0,1)^{d' \times Nd'}$. The following lemma allows gives us a lower bound on the gaussian width of each projection:

Lemma 4.3. *Let $\phi \sim \frac{1}{\sqrt{d'}}\mathcal{N}(0,1)^{d' \times Nd'}$. Then for $\widehat{\mathsf{Sp}}_\ell \subset \mathcal{S}^{Nd'-1}$,*

$$\mathop{\mathbf{E}}_{\phi}\left[\omega(\phi^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell))\right] \geqslant (1-4\varepsilon)\sqrt{d'} \tag{5}$$

From Lemma 4.3, we establish the following lower bound on the gaussian width of the projections of $\widehat{\mathsf{Sp}}_\ell^{Nd'-1}$:

$$\mathbb{E}_{\phi_i}\Big[\omega(\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell))\Big] \geqslant (1-4\varepsilon)\sqrt{d'} \tag{6}$$

Now, we argue that because $\phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)$ has gaussian width close to $\sqrt{d'}$, it in fact covers any fixed point on $\mathcal{S}^{d'-1}$ with high probability. We start by stating the following lemma on concentration of minimum distance w.r.t. large gaussian width sets.

Lemma 4.4. *Let $T \subset \mathcal{S}^{d'-1}$ be such that $\omega(T) \geqslant \sqrt{d'}(1-\varepsilon)$. Let $R : \mathbb{R}^{d'} \mapsto \mathbb{R}^{d'}$ be a uniform random rotation operator. Then for any $\mathbf{x} \in \mathcal{S}^{d'-1}$*

$$\mathop{\mathbf{Pr}}_{R}\left[\min_{\mathbf{y} \in R(T)} \|\mathbf{x}-\mathbf{y}\|_2 \geqslant 2\varepsilon^{1/4}\right] \leqslant \exp\Big(-O\Big(\frac{k}{\varepsilon}\log\frac{d}{k}\Big)\Big)$$

Observe that the distribution of ϕ_i^{norm} is rotation-invariant; see A.4 for a formal proof. Fixing $\mathbf{x} \in \mathcal{S}^{d'-1}$, we invoke Lemma 4.4 in our setting:

$$\mathop{\mathbf{Pr}}_{\phi_i}\Big(\min_{\mathbf{y} \in \phi_i^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)} \|\mathbf{y}-\mathbf{x}\|_2 > 16\varepsilon^{\frac{1}{4}}\Big) \leqslant \exp\Big(-O\Big(\frac{k}{\varepsilon}\log\frac{m}{k}\Big)\Big) \tag{7}$$

Let P_ε be an ε-cover of $\mathcal{S}^{d'-1}$ such that $|P_\varepsilon| = O(1/\varepsilon)^{d'}$. Then:

$$\begin{aligned}
&\mathop{\mathbf{Pr}}_{\Phi}\left(\exists \mathbf{x} \in P_\varepsilon : \min_{\mathbf{y}\in\Phi^{\mathrm{norm}}(\widehat{\mathsf{Sp}}_\ell)} \|\mathbf{y}-\mathbf{x}\|_2 \geqslant 16\varepsilon^{\frac{1}{4}}\right)\\
&\leqslant |P_\varepsilon| \mathop{\mathbf{Pr}}_{\Phi}\left(\forall \quad i \in [L] \min_{\mathbf{y}\in\phi_i^{\mathrm{norm}}(\widehat{\mathsf{Sp}}_\ell)} \|\mathbf{y}-\mathbf{x}\|_2 \geqslant 16\varepsilon^{\frac{1}{4}}\right)\\
&\overset{1}{=} |P_\varepsilon| \prod_{i=1}^{L} \mathop{\mathbf{Pr}}_{\phi_i}\left(\min_{\mathbf{y}\in\phi_i^{\mathrm{norm}}(\widehat{\mathsf{Sp}}_\ell)} \|\mathbf{y}-\mathbf{x}\|_2 \geqslant 16\varepsilon^{\frac{1}{4}}\right)\\
&= |P_\varepsilon| \left(\mathop{\mathbf{Pr}}_{\phi_i}\left(\min_{\mathbf{y}\in\phi_i^{\mathrm{norm}}(\widehat{\mathsf{Sp}}_\ell)} \|\mathbf{y}-\mathbf{x}\|_2 \geqslant 16\varepsilon^{\frac{1}{4}}\right)\right)^{L}\\
&\leqslant |P_\varepsilon| \exp\Big(-O\Big(\frac{k}{\varepsilon}\log\frac{m}{k}\Big)\Big)^{\varepsilon^2 O(\sqrt{m/k})}\\
&\leqslant \exp\Big(\Big(\log\frac{C'}{\varepsilon}\Big)\Big(\frac{k}{\varepsilon^2}\log\frac{m}{k}\Big) - k\varepsilon\sqrt{\frac{m}{k}}\log\frac{m}{k}\Big)\\
&\leqslant \exp\Big(-O\Big(k\log\frac{m}{k}\Big)\Big)
\end{aligned}$$

where step 1 uses the independence of the ϕ_i's and the last step uses the fact that $\varepsilon \geqslant \left(\frac{k}{m}\right)^{\frac{1}{8}}$. The above inequality states that $\Phi^{\mathrm{norm}}(\widehat{\mathsf{Sp}}_\ell)$ is an $16\varepsilon^{1/4}+\varepsilon$ cover of $\mathcal{S}^{d'-1}$ with positive probability, which completes the proof.

□

5 Conclusion

Our work opens the possibility of using linear queries to efficiently test other properties which arise in machine learning and convex optimization. Some questions directly motivated by this work are:

Other notions of distance: Our work uses isoangularity to measure how far the input vectors are from a sparse encoding. A relevant direction of future work is to complement this formulation by quantifying far in terms of other distance measures, such as the Hausdorff distance between two sets of vectors under the best possible alignment.

Faster testing of sparse recovery: One can study a special case of the dictionary testing problem in which the dictionary is known before hand. Note that if the given dictionary satisfies RIP, then one can use techniques from the compressed sensing literature to solve this problem exactly. Still, the question of whether we can devise more efficient algorithms just for testing that the input vectors admit a sparse representation (with respect to the given dictionary) remains an interesting open problem.

Other restrictions on the dictionary: Another important direction of future work is to consider the dictionary testing and sparse recovery problems in the context of commonly used dictionaries, such as the ones composed of Fourier basis, wavelet basis, and ridgelets. In particular, these dictionaries do satisfy RIP, but given their applicability it is relevant to understand if the results obtained in this paper can be strengthened with these additional restrictions on the dictionary.

Construction of ε-nets: A key technical contribution of the paper is to show that an ε-net of the unit sphere can be obtained by projecting down (from an appropriately larger dimension) the set of sparse vectors; see Lemma 4.2. It might be of independent interest to understand if one can obtain such nets by projecting other structured sets with high gaussian width.

References

[AAJ+14] Alekh Agarwal, Animashree Anandkumar, Prateek Jain, Praneeth Netrapalli, and Rashish Tandon. Learning sparsely used overcomplete dictionaries. In *Proc. 27th Annual ACM Workshop on Computational Learning Theory*, pages 123–137, 2014. 2, 5

[AEB05] Michal Aharon, Michael Elad, and Alfred M Bruckstein. K-svd and its non-negative variant for dictionary design. In *Optics & Photonics 2005*, pages 591411–591411. International Society for Optics and Photonics, 2005. 5

[AGM14] Sanjeev Arora, Rong Ge, and Ankur Moitra. New algorithms for learning incoherent and overcomplete dictionaries. In *Proc. 27th Annual ACM Workshop on Computational Learning Theory*, pages 779–806, 2014. 5

[Aha06] Michal Aharon. *Overcomplete dictionaries for sparse representation of signals*. PhD thesis, Technion - Israel Institute of Technology, 2006. 5

[BDKT12] Aditya Bhaskara, Daniel Dadush, Ravishankar Krishnaswamy, and Kunal Talwar. Unconditional differentially private mechanisms for linear queries. In *Proc. 44th Annual ACM Symposium on the Theory of Computing*, pages 1269–1284. ACM, 2012. 3

[BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, 1991. 1

[BKS15] Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proc. 47th Annual ACM Symposium on the Theory of Computing*, pages 143–151, New York, NY, USA, 2015. ACM. 5

[BLM13] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford, 2013. 6

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comp. Sys. Sci.*, 47:549–595, 1993. Earlier version in STOC'90. 1

[CRT06] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006. 2

[CSZ00] Artur Czumaj, Christian Sohler, and Martin Ziegler. Property testing in computational geometry. In *Proc. 8th European Symposium on Algorithms*, pages 155–166. Springer, 2000. 5

[DH01] David L Donoho and Xiaoming Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, 47(7):2845–2862, 2001. 2

[DMA97] Geoff Davis, Stephane Mallat, and Marco Avellaneda. Adaptive greedy approximations. *Constructive approximation*, 13(1):57–98, 1997. 5

[EA06] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006. 2

[EAH99] Kjersti Engan, Sven Ole Aase, and J Hakon Husoy. Method of optimal directions for frame design. In *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 2443–2446. IEEE, 1999. 5

[EP07] A Evgeniou and Massimiliano Pontil. Multi-task feature learning. In *Proc. 19th Advances in Neural Information Processing Systems*, page 41, 2007. 2

[GGR98] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45:653–750, 1998. 1

[GMS03] Anna C Gilbert, S Muthukrishnan, and Martin J Strauss. Approximation of functions over redundant dictionaries using coherence. In *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 243–252. Society for Industrial and Applied Mathematics, 2003. 2

[Gor85] Yehoram Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50(4):265–289, 1985. 5

[HKR12] Justin Hsu, Sanjeev Khanna, and Aaron Roth. Distributed private heavy hitters. In *Proc. 39th Annual International Conference on Automata, Languages, and Programming*, pages 461–472. Springer, 2012. 3

[HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proc. 42nd Annual ACM Symposium on the Theory of Computing*, pages 705–714. ACM, 2010. 3

[KS03] Robert Krauthgamer and Ori Sasson. Property testing of data dimensionality. In *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 18–27. Society for Industrial and Applied Mathematics, 2003. 5

[LS00] Michael S Lewicki and Terrence J Sejnowski. Learning overcomplete representations. *Neural Comp.*, 12(2):337–365, 2000. 2

[MLB+08] Julien Mairal, Marius Leordeanu, Francis Bach, Martial Hebert, and Jean Ponce. Discriminative sparse image models for class-specific edge detection and image interpretation. In *Proc. 10th European Conference on Computer Vision*, pages 43–56. Springer, 2008. 2

[MSE08] Julien Mairal, Guillermo Sapiro, and Michael Elad. Learning multiscale sparse representations for image and video restoration. *Multiscale Modeling & Simulation*, 7(1):214–241, 2008. 2

[OF96] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996. 2

[OF97] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997. 2, 5

[Ols02] Bruno A Olshausen. Sparse coding of time-varying natural images. *Journal of Vision*, 2(7):130–130, 2002. 2

[Pey09] Gabriel Peyré. Sparse modeling of textures. *Journal of Mathematical Imaging and Vision*, 34(1):17–31, 2009. 2

[RBL07] Marc'Aurelio Ranzato, Y-Lan Boureau, and Yann LeCun. Sparse feature learning for deep belief networks. In *Proc. 20th Advances in Neural Information Processing Systems*, pages 1185–1192, 2007. 2

[Ron08] Dana Ron. Property testing: A learning theory perspective. *Foundations and Trends in Machine Learning*, 1(3):307–402, 2008. 1

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. on Comput.*, 25:252–271, 1996. 1

[RS06] Ronitt Rubinfeld and Asaf Shapira. Sublinear time algorithms. In *Proc. International Congress of Mathematicians 2006*, volume 3, pages 1095–1110, 2006. 1

[RV08] Mark Rudelson and Roman Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008. 4

[Sle62] David Slepian. The one-sided barrier problem for gaussian noise. *The Bell System Technical Journal*, 41(2):463–501, 1962. 6

[SWW12] Daniel A Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. In *Proc. 25th Annual ACM Workshop on Computational Learning Theory*, pages 37–1, 2012. 5

[Ver15] Roman Vershynin. Estimation in high dimensions: a geometric perspective. In *Sampling Theory, a Renaissance*, pages 3–66. Springer, 2015. 4

[YWHM08] Jianchao Yang, John Wright, Thomas Huang, and Yi Ma. Image super-resolution as sparse representation of raw image patches. In *Proc. 26th IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2

A Supplementary Lemmas

A.1 Proof of Lemma 3.1

Proof. First we prove the upper bound. Let $\Psi : X \mapsto S$ be the ε-isometric embedding map. Given $\mathbf{g} \sim N(0, \sqrt{1+\varepsilon})^m$ and $\mathbf{h} \sim N(0,1)^d$, we define the gaussian processes $\{G_{\mathbf{x}}\}_{\mathbf{x}\in X}$ and $\{H_{\mathbf{x}}\}_{\mathbf{x}\in X}$ as follows

$$\begin{aligned} G_{\mathbf{x}} &\stackrel{\text{def}}{=} \mathbf{g}^\top \mathbf{x} \\ H_{\mathbf{x}} &\stackrel{\text{def}}{=} \mathbf{h}^\top \Psi(\mathbf{x}) \end{aligned}$$

Fix $\mathbf{x}, \mathbf{y} \in X$. Using the ε-isometry of Ψ we get:

$$\begin{aligned} \mathop{\mathbf{E}}_{\mathbf{h}\sim N(0,1)^d}\left[\left|H_{\mathbf{x}} - H_{\mathbf{y}}\right|^2\right] &= \mathop{\mathbf{E}}_{\mathbf{h}\sim N(0,1)^d}\left[\left|\mathbf{h}^\top\Psi(\mathbf{x}) - \mathbf{h}^\top\Psi(\mathbf{y})\right|^2\right] \\ &\stackrel{1}{=} \|\Psi(\mathbf{x}) - \Psi(\mathbf{y})\|^2 \\ &\stackrel{2}{\leqslant} (1+\varepsilon)\|\mathbf{x}-\mathbf{y}\|^2 \\ &= \mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,\sqrt{1+\varepsilon})^m}\left[\left|G_{\mathbf{x}} - G_{\mathbf{y}}\right|^2\right] \end{aligned}$$

where step 1 follows from the fact that the variance h in the direction of a vector $\mathbf{v}$ is $\|\mathbf{v}\|_2^2$, and inequality 2 follows from the isometric property. Therefore, using Lemma 2.5, we have

$$\mathop{\mathbf{E}}_{\mathbf{h}\sim N(0,1)^d}\left[\max_{\mathbf{x}\in S} \mathbf{g}^\top\Psi(\mathbf{x})\right] \leqslant \mathop{\mathbf{E}}_{\mathbf{g}\sim\sqrt{1+\varepsilon}N(0,1)^m}\left[\max_{\mathbf{x}\in S}\mathbf{g}^\top\mathbf{x}\right] \tag{8}$$

which directly gives us $\omega(S) \leqslant \sqrt{1+\varepsilon}\cdot\omega(X) \leqslant (1+\varepsilon)\omega(X)$. The other direction follows by using the lower bound given by isometry.

□

A.2 Proof of Lemma 4.3

The proof of the lemma proceeds in two steps: we first restrict to the case where the maximum $\|\cdot\|_2$ length of the projected vectors is not much larger than the expected value. This is done by using Gordon's theorem and Lipschitz concentration for gaussians (see Theorem 2.1 and Corollary 2.3). Following that, we observe that conditioning the expectation by this high probability event on the length of the projected vectors does not affect the expectation by much (Lemma A.1). The rest of the proof follows using standard estimates of gaussian widths of $\widehat{\mathsf{Sp}}_\ell$ (see Lemma A.3).

Upper Bound on the $\|\cdot\|_2$ length: By setting $D = Nd'$ in Theorem 2.1, we upper bound the expected maximum $\|\cdot\|_2$ length :

$$\mathop{\mathbf{E}}_{\phi}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\|\phi(\mathbf{x})\|_2\right] \leqslant 1 + \frac{\omega(\widehat{\mathsf{Sp}}_\ell)}{\sqrt{d'}} \tag{9}$$

Furthermore, from Lemma A.3, we have $\omega(\widehat{\mathsf{Sp}}_\ell) = \sqrt{C_0\ell\log\frac{Nd'}{\ell}}$ for some constant $C_0 > 0$. Therefore by our choice of parameters we have:

$$\frac{\omega(\widehat{\mathsf{Sp}}_\ell)}{\sqrt{d'}} = \Big(\frac{Ck}{\varepsilon^2}\log\frac{d}{k}\Big)^{-\frac{1}{2}}\sqrt{C_0\ell\log\frac{Nd'}{\ell}} = \frac{C'}{\varepsilon} \tag{10}$$

Note that Eq. 10 holds with an equality for some constant C'. Now consider the event $\mathcal{E}$ where the maximum $\|\cdot\|_2$ length is at most $1 + \frac{C'}{\varepsilon}(1+\varepsilon)$. Then by using gaussian concentration for Lipschitz functions (Corollary 2.3), we upper bound probability of the event $\neg\mathcal{E}$:

$$\begin{aligned}\mathop{\mathbf{Pr}}_{\phi\sim\frac{1}{\sqrt{d'}}N(0,1)^{d'\times Nd'}}\left(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\|\phi(\mathbf{x})\|_2 \geqslant 1+\frac{C'}{\varepsilon}(1+\varepsilon)\right) &\leqslant \exp\Big(-c\varepsilon^2\omega(\widehat{\mathsf{Sp}}_\ell)^2\Big) \qquad \text{(Using Corollary2.3)}\\ &\leqslant \exp\Big(-O\Big(\frac{k}{\varepsilon^2}\log\frac{d}{k}\Big)\Big) = \varepsilon_{k,d} \end{aligned}\tag{11}$$

where $\varepsilon_{k,d}$ can be made arbitrarily small by choosing d large enough.

Lower Bounding the gaussian width : Recall that the operator ϕ^{norm} is defined as $\phi(\mathbf{x}) \stackrel{\text{def}}{=} \phi(\mathbf{x})/\|\phi(\mathbf{x})\|_\mathbf{2}$. The operational expression for the gaussian width of the projected set restricted to coordinates in $[Nd']$ is given by :

$$\omega(\phi^{\text{norm}}(\widehat{\mathsf{Sp}}_\ell)) = \mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi^{\text{norm}}(\mathbf{x})\right] \tag{12}$$

We shall also need the following lemma which states that conditioning by the large probability event $\mathcal{E}$ does not reduce the expectation by much.

Lemma A.1. *There exists universal constants ℓ, d_0 such that for all $d \geqslant d_0$ and $d \geqslant k \geqslant \ell$ and the event $\mathcal{E}$ defined as above, we have*

$$\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'},\phi}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\Big|\mathcal{E}\right] \geqslant \mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'},\phi}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\right] - \gamma_{k,d}$$

where $\gamma_{k,d}$ decays exponentially in d

We defer the proof of the Lemma to Section A.3.1. Equipped with the above, we proceed to lower bound the expected Gaussian width:

$$
\begin{aligned}
&\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi^{\mathrm{norm}}(\mathbf{x})\right]\\
\geqslant&(1-\varepsilon_{k,d})\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi^{\mathrm{norm}}(\mathbf{x})\quad\middle|\quad\mathcal{E}\right]\\
\geqslant&\frac{(1-\varepsilon_{k,d})}{1+\frac{C'}{\varepsilon}(1+\varepsilon)}\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\quad\middle|\quad\mathcal{E}\right]\\
\overset{1}{\simeq}&\frac{(1-\varepsilon_{k,d})}{1+\frac{C'}{\varepsilon}(1+\varepsilon)}\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\right]
\end{aligned}
\tag{13}
$$

where the first inequality follows from the fact that $\mathcal{E}$ is a large probability event, and step 1 follows from Lemma A.1. Removing the conditioning allows us to relate the expectation term to the gaussian width of $\widehat{\mathsf{Sp}}_\ell$. Let B denote the event that $\|\mathbf{g}\|_2 \in [\sqrt{d'}(1-\varepsilon), \sqrt{d'}(1+\varepsilon)]$. Then,

$$
\begin{aligned}
\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\right]\geqslant&(1-\varepsilon)\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\middle|B\right]\\
\overset{2}{=}&(1-\varepsilon)\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\mathop{\mathbf{E}}_{\tilde{\mathbf{g}}\sim N\left(0,\frac{\|\mathbf{g}\|}{\sqrt{d'}}\right)^{Nd'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\tilde{\mathbf{g}}^\top\mathbf{x}\middle|B\right] &(14)\\
\overset{3}{\geqslant}&(1-\varepsilon)\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\mathop{\mathbf{E}}_{\tilde{\mathbf{g}}\sim N(0,1-\varepsilon)^{Nd'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\tilde{\mathbf{g}}^\top\mathbf{x}\middle|B\right]\\
=&(1-\varepsilon)\mathop{\mathbf{E}}_{\tilde{\mathbf{g}}\sim N(0,1-\varepsilon)^{Nd'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\tilde{\mathbf{g}}^\top\mathbf{x}\right]\\
\geqslant&(1-\varepsilon)^2\mathop{\mathbf{E}}_{\tilde{\mathbf{g}}\sim N(0,1)^{Nd'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\tilde{\mathbf{g}}^\top\mathbf{x}\right] &(15)\\
\geqslant&(1-\varepsilon)^2\omega(\widehat{\mathsf{Sp}}_\ell)
\end{aligned}
\tag{16}
$$

In step 2, in the inner expectation $\mathbf{g}\in\mathbb{R}^{d'}$ is a fixed vector, and therefore $\mathbf{g}^T\phi$ is distributionally equivalent to a gaussian vector in $\mathbb{R}^{Nd'}$, scaled by $\frac{\|\mathbf{g}\|_2}{\sqrt{d'}}$. Step 3 follows from the lower bound on the $\|\cdot\|_2$-length of $\mathbf{g}$. Plugging in the lower bound on the expectation term, we get :

$$
\begin{aligned}
\mathop{\mathbf{E}}_{\phi}\mathop{\mathbf{E}}_{\mathbf{g}\sim N(0,1)^{d'}}\left[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi^{\mathrm{norm}}(\mathbf{x})\right]&\geqslant(1-\varepsilon)^2\frac{(1-\varepsilon_{k,d})}{1+\frac{C'}{\varepsilon}(1+\varepsilon)}\omega(\widehat{\mathsf{Sp}}_\ell)\\
&=(1-\varepsilon)^2\frac{(1-\varepsilon_{k,d})}{1+\frac{C'}{\varepsilon}(1+\varepsilon)}\left(\frac{C'}{\varepsilon}\sqrt{d'}\right)\\
&\geqslant\sqrt{d'}(1-4\varepsilon)
\end{aligned}
$$

for sufficiently small ε and large d.

A.3 Proof of Lemma 4.4

We begin by looking at the expression of the square of the $\|\cdot\|_2$ distance. For any fixed $\mathbf{x}, \mathbf{y} \in \mathcal{S}^{d'-1}$, we have

$$\|\mathbf{y} - \mathbf{x}\|_2^2 = 2 - 2\mathbf{x}^\top \mathbf{y} \tag{17}$$

Therefore, minimizing the $\|\cdot\|_2$ norm would be equivalent to maximizing the dot product term. Furthermore, it is known that a random gaussian vector $\mathbf{g} \sim N(0,1)^{d'-1}$ can be rewritten as $Z\mathbf{r}$ where $\mathbf{r} \overset{\text{unif}}{\sim} \mathcal{S}^{d'-1}$ and Z is a χ^2-random variable with d' degrees of freedom. Using this decomposition, we get :

$$\begin{aligned}
&\Pr_{\mathbf{g}\sim N(0,1)^{d'}}\left(\max_{\mathbf{y}\in T} \mathbf{g}^\top \mathbf{y} \leqslant \sqrt{d'}(1-\varepsilon-\sqrt{\varepsilon})\right) \\
=&\Pr_{\mathbf{r},Z}\left(Z\max_{\mathbf{y}\in T} \mathbf{r}^\top \mathbf{y} \leqslant \sqrt{d'}(1-\varepsilon-\sqrt{\varepsilon})\right) \\
\geqslant&\Pr_{\mathbf{r},Z}\left(Z\max_{\mathbf{y}\in T} \mathbf{r}^\top \mathbf{y} \leqslant \sqrt{d'}(1-\varepsilon-\sqrt{\varepsilon}) \middle| Z \leqslant \sqrt{d'}(1+\varepsilon)\right)\Pr\left(Z \leqslant \sqrt{d'}(1+\varepsilon)\right) \\
\overset{1}{\geqslant}&\Pr_{\mathbf{r},Z}\left(Z\max_{\mathbf{y}\in T} \mathbf{r}^\top \mathbf{y} \leqslant \sqrt{d'}(1-\varepsilon-\sqrt{\varepsilon}) \middle| Z \leqslant \sqrt{d'}(1+\varepsilon)\right)(1-\varepsilon) \\
\geqslant&\frac{1}{2}\Pr_{\mathbf{r}}\left(\max_{\mathbf{y}\in T} \mathbf{r}^\top \mathbf{y} \leqslant \frac{1-\varepsilon-\sqrt{\varepsilon}}{1+\varepsilon}\right) \\
\geqslant&\frac{1}{2}\Pr_{\mathbf{r}}\left(\max_{\mathbf{y}\in T} \mathbf{r}^\top \mathbf{y} \leqslant (1-2\sqrt{\varepsilon})\right)
\end{aligned}$$

where in step 1, we used concentration for χ^2-random variables. We now relate the behavior of the maximum dot product of a set w.r.t a random vector to the maximum dot product of a fixed vector w.r.t. a randomly rotated set.

$$\begin{aligned}
\Pr_{\mathbf{r}}\left(\max_{\mathbf{y}\in T} \mathbf{r}^\top \mathbf{y} \leqslant (1-2\sqrt{\varepsilon})\right) &\overset{2}{=} \Pr_{R}\left(\max_{\mathbf{y}\in T} R(\mathbf{x})^\top \mathbf{y} \leqslant (1-2\sqrt{\varepsilon})\right) \\
&= \Pr_{R}\left(\max_{\mathbf{y}\in T} \mathbf{x}^\top R^{-1}(\mathbf{y}) \leqslant (1-2\sqrt{\varepsilon})\right) \\
&\overset{3}{=} \Pr_{R}\left(\max_{\mathbf{y}\in T} \mathbf{x}^\top R(\mathbf{y}) \leqslant (1-2\sqrt{\varepsilon})\right) \\
&= \Pr_{R}\left(\max_{\mathbf{y}\in R(T)} \mathbf{x}^\top \mathbf{y} \leqslant (1-2\sqrt{\varepsilon})\right)
\end{aligned} \tag{18}$$

Step 2 follows from the fact that applying a uniformly random rotation on a unit vector is equivalent to sampling u.a.r. from the unit sphere $\mathcal{S}^{d'-1}$, and step 3 follows from the fact that if R is u.a.r. rotation, then R^{-1} is also a u.a.r rotation. Furthermore, using gaussian concentration for Lipschitz functions :

$$\Pr_{\mathbf{g}\sim N(0,1)^{d'}}\left(\max_{\mathbf{y}\in T} \mathbf{g}^\top \mathbf{y} \leqslant \sqrt{d'}(1-\varepsilon-\sqrt{\varepsilon})\right) \leqslant \exp\left(-O\Big(\frac{k}{\varepsilon}\log\frac{m}{k}\Big)\right) \tag{19}$$

and the l.h.s is a lower bound on $\frac{1}{2}\Pr_R\left(\max_{\mathbf{y}\in R(T)} \mathbf{x}^\top \mathbf{y} \leqslant (1-2\sqrt{\varepsilon})\right)$ (Eq. 18). Therefore rearranging the equations, we have:

$$\Pr_{R}\left(\min_{\mathbf{y}\in R(T)} \|\mathbf{x}-\mathbf{y}\| \geqslant 2\varepsilon^{1/4}\right) = \Pr_{R}\left(\max_{\mathbf{y}\in R(T)} \mathbf{x}^\top \mathbf{y} \leqslant (1-2\sqrt{\varepsilon})\right) \leqslant 2\exp\left(-O\Big(\frac{k}{\varepsilon}\log\frac{m}{k}\Big)\right) \tag{20}$$

A.3.1 Proof of Lemma A.1

Proof. The proof uses the more general observation relating conditional expectations to their unconditioned counterparts :

Proposition A.2. *Let $\mathcal{E}$ be an event such that $\mathbf{Pr}(\mathcal{E}) \geqslant 1-\eta$. Let $Z : \mathbb{R}^{d'} \mapsto \mathbb{R}_+$ be a non-negative random variable. Let $t_0 > 0$ be such that*

$$\int_{t_0}^{\infty} t\,\mathbf{Pr}(Z \geqslant t)dt \leqslant \alpha$$

Then, the following holds true : $\mathbf{E}[Z|\mathcal{E}] \geqslant \mathbf{E}[Z] - \frac{\eta t_0^2}{2} - \alpha$.

Proof. We begin by observing that for any event B,

$$\mathbf{Pr}(B|\mathcal{E}) \geqslant \mathbf{Pr}(B) - \eta$$

which follows from the definition of conditional expectation. Therefore,

$$\begin{aligned}
\mathbf{E}[Z|\mathcal{E}] = \int_0^{\infty} t\,\mathbf{Pr}(Z \geqslant t|\mathcal{E})dt &\geqslant \int_0^{t_0} t\,\mathbf{Pr}(Z \geqslant t|\mathcal{E})dt \\
&\geqslant \int_0^{t_0} t\big(\mathbf{Pr}(Z \geqslant t) - \eta\big)dt \\
&\geqslant \int_0^{t_0} \mathbf{Pr}(Z \geqslant t)tdt - \int_0^{t_0} t\eta dt \\
&= \mathbf{E}[Z] - \alpha - \frac{\eta t_0^2}{2}
\end{aligned}$$

□

We apply the above lemma to our setting: let $\mathcal{E}$ be the event as described in the proof of lemma 4.3 and let Z be the random variable :

$$Z = \max_{\mathbf{x} \in \widehat{\mathsf{Sp}}_\ell} \mathbf{g}^\top \phi(\mathbf{x})$$

From Eq. 11, we know that $\mathbf{Pr}(\neg\mathcal{E}) = \eta = \varepsilon_{k,d}$. Abusing notation, we denote $\omega^* = \omega(\widehat{\mathsf{Sp}}_\ell)$. Let $t_0 = 4\omega(\widehat{\mathsf{Sp}}_\ell) = 4\omega^*$. For a fixed choice of $\delta > 0$, let B_δ be the event that $\|\mathbf{g}\|_2 \leqslant \sqrt{d'}(2+\delta)$. Then,

$$\begin{aligned}
&\mathop{\mathbf{Pr}}_{\mathbf{g}\sim N(0,1)^{d'},\phi}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell} \mathbf{g}^\top\phi(\mathbf{x}) \geqslant (2+\delta)^2\omega^*\Big) \\
&\leqslant \mathbf{Pr}(\neg B_\delta) + \mathop{\mathbf{Pr}}_{\mathbf{g}\sim N(0,1)^{d'},\phi}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell} \mathbf{g}^\top\phi(\mathbf{x}) \geqslant (2+\delta)^2\omega^* \Big| B_\delta\Big) \\
&\overset{1}{\leqslant} \exp\Big(-O((1+\delta)^2 d')\Big) + \mathop{\mathbf{Pr}}_{\mathbf{g}\sim N(0,1)^{d'},\phi}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell} \mathbf{g}^\top\phi(\mathbf{x}) \geqslant (2+\delta)^2\omega^* | B_\delta\Big) \qquad (21) \\
& \qquad (22)
\end{aligned}$$

where inequality 1 follows by concentration on χ^2 variables. We upper bound the remaining probability term as :

$$
\begin{aligned}
&\underset{\mathbf{g}\sim N(0,1)^{d'},\phi}{\mathbf{Pr}}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\geqslant(2+\delta)^2\omega^*|B_\delta\Big)\\
&\overset{2}{\leqslant}\underset{\tilde{\mathbf{g}}\sim N(0,1)^{Nd'}}{\mathbf{Pr}}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\tilde{\mathbf{g}}^\top\mathbf{x}\geqslant\frac{(2+\delta)^2}{1+\delta}\omega^*|B_\delta\Big)\\
&\leqslant\underset{\tilde{\mathbf{g}}\sim N(0,1)^{Nd'}}{\mathbf{Pr}}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\tilde{\mathbf{g}}^\top\mathbf{x}\geqslant(2+\delta)\omega^*\Big) \qquad (23)\\
&\overset{3}{\leqslant}\exp\Big(-O((1+\delta)\omega^*)^2\Big)\\
&\leqslant\exp\Big(-O\Big((1+\delta)^2k\log\frac{d}{k}\Big)\Big) \qquad (24)
\end{aligned}
$$

Step 2 can be shown using arguments identical to the ones used in steps 14-15 in the proof of lemma 3.2, and step 3 follows from gaussian concentration. Now we proceed to upper bound the quantity α (as in Proposition A.2):

$$
\begin{aligned}
\int_{t_0}^{\infty}t\,\mathbf{Pr}(Z\geqslant t)dt &= \int_{t_0}^{\infty}(4\omega^*+t)\underset{\mathbf{g}\sim N(0,1)^{d'},\phi}{\mathbf{Pr}}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\geqslant 4\omega^*+t\Big)dt\\
&\overset{1}{=}\int_0^{\infty}2(2+\delta)^3(\omega^*)^2\underset{\mathbf{g}\sim N(0,1)^{d'},\phi}{\mathbf{Pr}}\Big(\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_\ell}\mathbf{g}^\top\phi(\mathbf{x})\geqslant(2+\delta)^3(\omega^*)^2\Big)d\delta\\
&= \int_0^{\infty}(2+\delta)^3(\omega^*)^2\exp\Big(-O((1+\delta)^2k\log\frac{d}{k})\Big)d\delta\\
&\leqslant \int_0^{\infty}\exp\Big(-O((1+\delta)^2k\log\frac{d}{k})\Big)d\delta=\beta_{k,d} \qquad (25)
\end{aligned}
$$

where step 1 is a change of variables argument where we set $t=(2+\delta)^2\omega^*-4\omega^*$. Plugging in the upper bounds from Equations 21,24 and 25, we get

$$\mathbf{E}[Z|\mathcal{E}]\geqslant\mathbf{E}[Z]-\alpha_{k,d}-\beta_{k,d}=\mathbf{E}[Z]-\gamma_{k,d}$$

where $\gamma_{k,d}$ decays exponentially in k.

□

A.3.2 Gaussian Width of the discretized sparse set $\widehat{\mathsf{Sp}}_k$

Lemma A.3. *Let $\widehat{\mathsf{Sp}}_\ell\subset\mathcal{S}^{m-1}$ be the discrete k-sparse set on the unit sphere. Then,*

$$\omega(\widehat{\mathsf{Sp}}_\ell)=\Theta\Big(\sqrt{\ell\log\frac{m}{\ell}}\Big) \qquad (26)$$

Proof. For the upper bound, observe that $\widehat{\mathsf{Sp}}_\ell\subset\mathsf{Sp}_\ell^m$ and gaussian width is monotonic. Therefore, from Lemma 1.6, we have $\omega(\widehat{\mathsf{Sp}}_k)=O\Big(\sqrt{k\log\frac{d}{k}}\Big)$. Towards proving the asymptotic lower bound: Given independent gaussians $g_1,\ldots,g_n\sim N(0,1)$, it is known that

$$\underset{g_1,\ldots,g_n}{\mathbf{E}}\Big[\max_{g_i}g_i\Big]\geqslant C_0\sqrt{\log n} \qquad (27)$$

for some constant $C_0>0$ independent of the number of gaussians. Now, without loss of generality let d be divisible by k. We partition the m coordinates into ℓ blocks $B_1,\ldots,B_\ell$ of $\frac{m}{\ell}$ coordinates each. Then,

$$\mathop{\mathbf{E}}_{\mathbf{g}}\Big[\max_{\mathbf{x}\in\widehat{\mathsf{Sp}}_k}\mathbf{g}^\top\mathbf{x}\Big] \geqslant \frac{1}{\sqrt{k}}\,\mathbf{E}\Big[\sum_{j\in[\ell]}\max_{g_{i_j}\in B_j} g_{i_j}\Big] \tag{28}$$

The inequality follows from the following observation: For any fixed realization of $\mathbf{g}\sim N(0,1)^m$, let i_j be the index of the maximum in the j^{th} block. Then there exists a vector in $\widehat{\mathsf{Sp}}_\ell$ which is supported on $i_1,\ldots,i_k$. Therefore, the dot product would be at least the sum of maximum Gaussians from each of the blocks scaled by $\frac{1}{\sqrt{\ell}}$. The lemma now follows from applying the lower bound from Eq. 27.

□

A.4 Rotational Invariance of ϕ^{norm}

Lemma A.4. *For any finite set $T\subset\mathcal{S}^{Nd'-1}$, the distribution of $\Phi^{\text{norm}}(S)$ is rotation invariant.*

Proof. Fix any $N=|T|$ vectors $\{\mathbf{z}'_1,\ldots,\mathbf{z}'_N\}$. Let $R:\mathbb{R}^{Nd'}\mapsto\mathbb{R}^{Nd'}$ be a fixed rotation. Then,

$$\begin{aligned}
\mathop{\mathbf{Pr}}_{\phi}\Big(\phi^{\text{norm}}(T)=\Big\{\frac{\mathbf{z}'_1}{\|\mathbf{z}'_1\|_2},\ldots,\frac{\mathbf{z}'_N}{\|\mathbf{z}'_N\|_2}\Big\}\Big) &= \mathop{\mathbf{Pr}}_{\phi}\Big(\phi(T)=\{\mathbf{z}'_1,\ldots,\mathbf{z}'_N\}\Big)\\
&= \mathop{\mathbf{Pr}}_{\phi}\Big(\phi(T)=\{R(\mathbf{z}'_1),\ldots,R(\mathbf{z}'_N)\}\Big)\\
&\overset{1}{=} \mathop{\mathbf{Pr}}_{\phi}\Big(\phi^{\text{norm}}(T)=\Big\{\frac{R(\mathbf{z}'_1)}{\|R(\mathbf{z}'_1)\|_2},\ldots,\frac{R(\mathbf{z}'_N)}{\|R(\mathbf{z}'_N)\|_2}\Big\}\Big)\\
&\overset{2}{=} \mathop{\mathbf{Pr}}_{\phi}\Big(\phi^{\text{norm}}(T)=\Big\{\frac{R(\mathbf{z}'_1)}{\|(\mathbf{z}'_1)\|_2},\ldots,\frac{R(\mathbf{z}'_N)}{\|(\mathbf{z}'_N)\|_2}\Big\}\Big)\\
&= \mathop{\mathbf{Pr}}_{\phi}\Big(\phi^{\text{norm}}(T)=R\Big(\Big\{\frac{\mathbf{z}'_1}{\|\mathbf{z}'_1\|_2},\ldots,\frac{\mathbf{z}'_N}{\|\mathbf{z}'_N\|_2}\Big\}\Big)\Big)
\end{aligned}$$

where step 1 follows from the rotational invariance of ϕ, and step 2 uses the observation that rotating a vector does not change it's ℓ_2-length. Since the equality holds for any rotation, the statement follows. □

A.4.1 From ε-covers to Isoangularity

Proposition A.5. *Let $S,\widehat{\mathsf{Sp}}\subset\mathcal{S}^{d-1}$ and $\tilde{S}\subset\mathbb{R}^d$ be sets of equal sizes. If*

1 *$\tilde{S}$ is ε-isometric to S and*

2 *$\widehat{\mathsf{Sp}}$ is an $O(\varepsilon^{1/4})$-cover of $\tilde{S}$*

then $\widehat{\mathsf{Sp}}$ is $O(\varepsilon^{1/8})$-isoangular to S.

Proof. Fix $\mathbf{x},\mathbf{y}\in S$. Then there exist $\tilde{\mathbf{x}},\tilde{\mathbf{y}}\in\tilde{S}$ (under the bijection of the isometry) and $\hat{\mathbf{x}},\hat{\mathbf{y}}\in\widehat{\mathsf{Sp}}$ (which are $C'\varepsilon^{1/4}$-close to $\tilde{\mathbf{x}},\tilde{\mathbf{y}}$ respectively). Since $S,\widehat{\mathsf{Sp}}$ are on the unit sphere, it is sufficient to show that $|\mathbf{x}^\top\mathbf{y}-\hat{\mathbf{x}}^\top\hat{\mathbf{y}}|=O(\varepsilon^{1/8})$. We prove this for near and faraway points separately.

For $O(\varepsilon^{1/4})$-Close points: Consider the case where $\|\mathbf{x}-\mathbf{y}\|\leqslant c_0\varepsilon^{1/4}$. Then,

$$\begin{aligned}
\|\hat{\mathbf{x}}-\hat{\mathbf{y}}\|_2 &\leqslant 2c_1\varepsilon^{1/4}+\|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2\\
&\leqslant 2c_1\varepsilon^{1/4}+\sqrt{1+\varepsilon}\|\mathbf{x}-\mathbf{y}\|_2\\
&\leqslant 2(c_0+c_1)\varepsilon^{1/4}=c'\varepsilon^{1/4}
\end{aligned}$$

Therefore,

$$\begin{aligned} c_0^2\varepsilon^{1/8} &\leqslant \|\mathbf{x}-\mathbf{y}\|_2^2 = 2 - 2\mathbf{x}^\top\mathbf{y} \\ \Rightarrow \mathbf{x}^\top\mathbf{y} &\geqslant 1 - c_0^2\varepsilon^{1/8} \end{aligned}$$

and $\mathbf{x}^\top\mathbf{y} \leqslant 1$ by Cauchy-Schwartz. Therefore, $\mathbf{x}^\top\mathbf{y} \in [1 - O(\varepsilon^{1/8}), 1]$, and by similar arguments we have $\hat{\mathbf{x}}^\top\hat{\mathbf{y}} \in [1 - O(\varepsilon^{1/8}), 1]$, which together gives us $|\mathbf{x}^\top\mathbf{y} - \hat{\mathbf{x}}^\top\hat{\mathbf{y}}| = O(\varepsilon^{1/8})$.

For faraway points: In this case we have $\|\mathbf{x}-\mathbf{y}\|_2 > c_0\varepsilon^{1/4}$. Then

$$\|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2 \geqslant c_0\varepsilon^{1/4}\sqrt{1-\varepsilon} \geqslant 0.5c_0\varepsilon^{1/4}$$

Therefore,

$$\begin{aligned} \|\hat{\mathbf{x}}-\hat{\mathbf{y}}\|_2^2 &\in \|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2^2 + c_1^2\varepsilon^{1/2} \pm \Big(2c_1\varepsilon^{1/4}\|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2\Big) \\ \Rightarrow \|\hat{\mathbf{x}}-\hat{\mathbf{y}}\|_2^2 &\in \|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2^2 + c_1^2\varepsilon^{1/2} \pm \Big(\frac{4c_1\varepsilon^{1/8}}{c_0}\|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2^2\Big) \\ \Rightarrow \|\hat{\mathbf{x}}-\hat{\mathbf{y}}\|_2^2 &\in \|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2^2 \pm \varepsilon^{1/8}\Big(\Big(4\frac{c_1}{c_0} + \frac{c_0^2}{c_1^2}\Big)\|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2^2\Big) \\ \Rightarrow \|\hat{\mathbf{x}}-\hat{\mathbf{y}}\|_2^2 &\in (1 \pm O(\varepsilon^{1/8}))\|\tilde{\mathbf{x}}-\tilde{\mathbf{y}}\|_2^2 \\ \Rightarrow \|\hat{\mathbf{x}}-\hat{\mathbf{y}}\|_2^2 &\in (1 \pm O(\varepsilon^{1/8}))(1\pm\varepsilon)\|\mathbf{x}-\mathbf{y}\|_2^2 \\ \Rightarrow 2 - 2\hat{\mathbf{x}}^\top\hat{\mathbf{y}} &\in (1 \pm O(\varepsilon^{1/8}))(2 - \mathbf{x}^\top\mathbf{y}) \\ \Rightarrow |\hat{\mathbf{x}}^\top\hat{\mathbf{y}} - \mathbf{x}^\top\mathbf{y}| &\leqslant O(\varepsilon^{1/8})(1 + \mathbf{x}^\top\mathbf{y}) = O(\varepsilon^{1/8}) \end{aligned}$$

which completes the proof for faraway points. □